

81-1708

ODP-81-668  
22 May 1981

Training 33

MEMORANDUM FOR: Chairman, Publications Review Board

THROUGH : Chief, Engineering Division  
Deputy Director for Processing, ODP  
Director of Data Processing  
Deputy Director for AdministrationSTAT FROM : [REDACTED]  
Engineering Division, ODP

SUBJECT : Request to Give a Presentation

1. I request permission to give a presentation describing Agency experience predicting the growth of computing needs.

2. When approved, I intend to speak at the [REDACTED] Conference in [REDACTED] during the week of August 23rd. The audience is expected to be comprised of about 400 persons from the United States and Canada representing organizations running similar computer systems.

STAT

3. None of the material to be presented is classified or controversial. I will describe methodologies that we have developed to anticipate the growth of interactive computing. Techniques and tools related to the IBM's VM/CMS timesharing system will be described in detail.

4. I am not under cover and will be identified as an Agency employee. I will also give the standard disclaimer that the views expressed are my own and not necessarily those of the Agency.

STAT



SUBJECT: Request to Give a Presentation

AUTHOR'S NAME:

STAT

TITLE OF PRESENTATION: VM Capacity Planning

I have reviewed the outline in paragraph 3 of this request, to the best of my knowledge have found it to be unclassified, and approve it for presentation.

**/s/ Bruce T. Johnson**

Bruce T. Johnson, D/ODP

13 AUG 1981

Date

**/s/ William N. Hart**

for Harry E. Fitzwater, DDA

14 AUG 1981

Date

Capacity Planning for a CMS Intensive Environment



STAT

Central Intelligence Agency  
Office of Data Processing  
Washington, D.C. 20505

Installation Code: CAD

VM/370 PROJECT

B583

ABSTRACT

The speaker will outline and discuss the capacity planning function at this large VM installation. Topics will include availability, establishing performance goals, tracking workload characteristics, data collection and reporting techniques, forecasting techniques, and cost-performance trade-offs. Many of the topics discussed can be applied at any size VM installation.

INTRODUCTION

A well-established capacity planning function has become a very valuable asset to the community of computer installations. In a time of rapidly increasing technology and decreasing prices, the ability to investigate all alternatives prior to procurement will provide that sought after cost-performance option. Effective capacity planning requires in-depth knowledge of the system and its user community. The items discussed in this paper will aide in establishing a new capacity planning function or improving an existing one.

AVAILABILITY

Service availability is the most important aspect of capacity planning. The impact of an outage is directly proportional to the size of the system. Most CMS intensive installations are prime shift oriented and any outage is considerable. The outcome of service outages with the most far reaching implications is the effect on user attitude. When the user population is unsure of the integrity of their data, they take certain precautions. For example, the user community might intermittently save the files they are editing to avoid losing and redoing work. This type of user reaction would have a significant effect on system activity.

At our installation, service availability is considered a very serious matter and the office has devoted considerable time and effort to improvements.

I assisted in the effort by developing a model to serve as a tracking aide and an analysis tool. I began by documenting each outage of a few selected services, then grouped these outages into what I termed 'subsystems'. A subsystem is defined as any system component capable of creating an outage (e.g. software, processor, channel, disk, power, operator,...). This provided the ability to locate the 'weakest link' in a service's availability. In order to compare all our services, I grouped the subsystems of each service into a standard set of categories. This allowed us to determine if one service was being unduly affected by a particular category. We could then investigate the subsystems of that service's category and locate the area of needed improvement. Attachment A is an example of the output of the model which has become a standardized report format at our installation. The sample report is for our main VM service and the category of USER SOFTWARE is marked (N/A) which denotes not applicable for this service. This category is for services supported by customer supplied software that executes under the operating system software. Since the model input is totally flexible, it also serves as a tool for the capacity planning analyst in predicting the availability of future configurations.

#### PERFORMANCE GOALS, DATA COLLECTION, AND REPORTING TECHNIQUES

Developing performance goals for a service must be accomplished in a reasonable manner with management's approval. Making a commitment to the user community requires the analysis of a well-established performance database. The IBM program product VMAP (5798-CPX) provides the tools required to develop such a database. The performance goals of most CMS intensive installations are oriented around trivial response or the truly interactive user, with the background users receiving the remaining system resources.

Once a reasonable and obtainable set of performance goals has been established, a concise report covering the service's performance, workload, and availability should be developed. Obtaining inputs from the planned recipients will provide a good outline. The report should also provide space for notes to explain poor performance and/or availability, their causes, and effects on workload. Also, the report should be produced weekly and contain previous weeks for comparison. Attachment B is an example of the report we developed.

#### WORKLOAD CHARACTERIZATION

A workload characterization is a definition of how your system is used. It is very valuable tool when studying performance and planning future systems. A workload characterization can be developed using several methods, but use the one best suited to your environment. The best method in a CMS intensive environment is command measurement. The information obtained from a command measurement study will provide the basis for a system profile. This profile will define how the users present the load to the system and provide input to the development of a benchmark system. It can also be used to locate high activity facilities which should be investigated for possible performance improvement. To obtain a complete system profile however, you should use other methods too (e.g. resource per user per second, logged-active user ratios,...). A workload characterization, like performance goals, should be developed from a well-established database.

## FORECASTING TECHNIQUES AND TREND ANALYSIS

A historical database is a mandatory requirement for forecasting and trend analysis. A database comprised of more than two years of data will allow analysis to avoid seasonal loads and plateaus. A good source for this database is tracking the resources per user required to maintain your performance goals. Another source is variables or measures that correlate well (.80 and up) with increases in system activity.

Another important component of forecasting, often difficult to discover or locate, is artificial constraints. Resource limitations not readily visible but present. For example, a high logon-logoff rate might reflect a deficient number of terminals. As mentioned before, availability could be a hidden constraint.

Forecasting is not just predicting user growth and justifying faster processors. It also involves predicting the effect of additional system resources (e.g. memory, paging devices, channels,...). However, when attempting any type of forecasting, it is extremely important to document the results, and state and explain all inputs and assumptions.

## COST-PERFORMANCE TRADE-OFFS

Probably the biggest fear in the computer industry today is arriving at the position of owing more on a piece of hardware than it's worth. Advancing technology and decreasing prices can make this a difficult task. Studying the system profile and other workload characterizations to determine the exact resource constraints will promote effective cost-performance analysis. The ability to define and evaluate equipment prior to procurement will allow you to determine exactly what is required.

Perhaps the best method of equipment evaluation is benchmarking. A benchmark system constructed from your system profile and other workload characterizations will allow evaluation of the equipment in your environment. However, there are many trade-offs between the complexity and accuracy of a benchmark system.

Cost-performance analysis does not just involve processors, it also includes other system components (e.g. DASD, memory,...). In the past, the area of VM DASD was fairly rigid, providing only 800 byte blocksizes. The following two charts show the value of the new file system's multiple blocksizes.

DASD Chart - percent utilization by blksize

Device	800	1K	2K	4K
3330-1	85.76	86.25	94.09	94.09
3350	79.57	80.40	85.76	85.76
3380	60.62	66.81	77.59	86.21

---

---

DASD Chart - blocks per 1 cylinder by blksize

---

Device	800	1K	2K	4K
3330-1	266	209	114	57
3350	570	450	240	120
3380	540	465	270	150

---

Notice that the 800 byte blocksize can present a problem when converting from 3350 to 3380 DASD in that the cylinders have less capacity, but in 1K blocksize format a 3380 cylinder has more capacity. Before converting DASD, it is a good idea to investigate the directory to determine user minidisk sizes. If the majority of your user minidisks are one cylinder, converting to DASD with cylinders twice the size (e.g. 3330 to 3350) may result in a higher total cost.

#### SUMMARY

Capacity planning is an important and necessary function in any computer installation. The ability to evaluate the capability of current and future system configurations is invaluable. A competent capacity planning function must consider all service components to determine the exact areas requiring improvement. (Remember: Service availability is the most critical component.) In order to operate properly and efficiently, the capacity planning function must have the blessing, attention, and respect of management.

## ATTACHMENT A: SAMPLE AVAILABILITY MODEL REPORT

VM/370 SYSTEM AVAILABILITY (01 APR - 30 JUN 1979)  
 SCHEDULED UP TIME: MON-FRI 0700-1800  
 REPORT BY SUBSYSTEM

	SUBSYSTEM	AVAILABILITY	MTBF	MDT
1	IBM 3033 CPU	99.638	116.908	0.425
2	MEMORY	100.000	0.000	0.000
3	DIRECTOR	99.908	703.350	0.650
4	CHANNEL - BYTE	100.000	0.000	0.000
5	CHANNEL - BLOCK	99.988	703.917	0.083
6	3033 CONSOLE	100.000	0.000	0.000
7	VM SOFTWARE	99.304	49.936	0.350
8	USER SOFT (N/A)	100.000	0.000	0.000
9	OTHER DASD	99.929	703.500	0.500
10	JES3	100.000	0.000	0.000
11	TBAR	99.948	351.817	0.183
12	IBM DRUM CONTR	100.000	0.000	0.000
13	IBM DRUM DASD	100.000	0.000	0.000
14	TELEX 3350 CONTR	99.373	77.731	0.491
15	TELEX 3350 DASD	99.212	87.306	0.694
16	COMTEN	99.983	703.883	0.117
17	IBM TAPE CONTR	100.000	0.000	0.000
18	IBM TAPE DRV	100.000	0.000	0.000
19	POWER	99.512	700.567	3.433
20	HARDWARE CHANGE (ED)	100.000	0.000	0.000
21	PROCEDURAL (OD)	100.000	0.000	0.000
22	UNKNOWN	100.000	0.000	0.000

SYSTEM - AVAIL: 96.84 MTBF: 15.723 MDT: 0.513

VM/370 SYSTEM AVAILABILITY (01 APR - 30 JUN 1979)  
 SCHEDULED UP TIME: MON-FRI 0700-1800  
 REPORT BY CATEGORY


	CATEGORY	AVAILABILITY	SUBSYSTEMS	MTBF	MDT
1	PROCESSOR	99.534	1- 6	87.590	0.410
2	O.S. SOFTWARE	99.304	7- 7	49.936	0.350
3	USER SOFTWARE (N/A)	100.000	8- 8	0.000	0.000
4	NETWORK	99.877	9-11	234.378	0.289
5	SYSTEM AND DB DASD	98.584	12-15	40.825	0.586
6	COMMUNICATIONS	99.983	16-16	703.883	0.117
7	TAPE SUBSYSTEM	100.000	17-18	0.000	0.000
8	ENVIRONMENT	99.512	19-19	700.567	3.433
9	HARDWARE MAINTENANCE	100.000	20-20	0.000	0.000
10	PROCEDURAL	100.000	21-21	0.000	0.000
11	UNKNOWN	100.000	22-22	0.000	0.000

SYSTEM - AVAIL: 96.84 MTBF: 15.723 MDT: 0.513


ATTACHMENT B: SAMPLE WEEKLY REPORT

WEEK : 43  
YEAR : FY-81VM TREND REPORT  
FOR LATEST TWENTY WORKING DAYS  
WEEK ENDING 24 JULY 81

## VM PERFORMANCE TRENDS



## VM WORKLOAD TRENDS



DATE : 07/03 07/10 07/17 07/24  
 AVAILABILITY : 99.85 92.88 99.64 98.58

## COMMENTS :

NOTES TO EXPLAIN PERFORMANCE PROBLEMS, WORKLOAD ENVIRONMENT,  
AND AVAILABILITY - SERVICE OUTAGES.